

Transformations Between Different Types of Unranked Bottom-Up Tree Automata

Xiaoxue Piao

Kai Salomaa

School of Computing, Queen's University
Kingston, Ontario K7L 3N6, Canada
{piao, ksalomaa}@cs.queensu.ca

We consider the representational state complexity of unranked tree automata. The bottom-up computation of an unranked tree automaton may be either deterministic or nondeterministic, and further variants arise depending on whether the horizontal string languages defining the transitions are represented by a DFA or an NFA. Also, we consider for unranked tree automata the alternative syntactic definition of determinism introduced by Cristau et al. (FCT'05, Lect. Notes Comput. Sci. 3623, pp. 68–79). We establish upper and lower bounds for the state complexity of conversions between different types of unranked tree automata.

Keywords: tree automata, unranked trees, state complexity, nondeterminism

1 Introduction

Descriptional complexity, or state complexity, of finite automata has been extensively studied in recent years, see [8, 10, 16, 18] and references listed there. On the other hand, very few papers explicitly discuss state complexity of tree automata. For classical tree automaton models operating on ranked trees [6, 7] many state complexity results are similar to corresponding results on string automata. For example, it is well known that determinizing an n state nondeterministic bottom-up tree automaton gives an automaton with at most 2^n states.

Modern applications of tree automata, such as XML document processing [12, 17], use automata operating on unranked trees. One approach is to first encode the unranked trees as binary trees [3]. The other approach that we consider here is to define the computation of the tree automaton directly on unranked XML-trees [2, 6, 17]. The set of transitions of an unranked tree automaton is, in general, infinite and the transitions are usually specified in terms of a regular language. Thus, in addition to the finite set of states used in the bottom-up computation, an unranked tree automaton needs for each state q and input symbol σ a finite string automaton to recognize the *horizontal language* consisting of strings of states defining the transitions associated to q and σ .

Here we consider bottom-up (frontier-to-root) unranked tree automata. Roughly speaking, we get different models depending on whether the bottom-up computation is nondeterministic or deterministic and whether the horizontal languages are recognized by an NFA or a DFA ((non-)deterministic finite automaton). Furthermore, there is more than one way to define determinism for unranked tree automata and we compare here two of the variants.

The more common definition [6, 17] requires that for any input symbol σ and two distinct states q_1, q_2 , the horizontal languages associated, respectively, with q_1 and σ and with q_2 and σ are disjoint. The condition guarantees that the bottom-up computation assigns a unique state to each node. To distinguish this from the syntactic definition of determinism of [5, 15], we call a deterministic tree automaton where the horizontal languages defining the transitions are specified by DFAs, a *weakly deterministic*

tree automaton. Note that a computation of a weakly deterministic automaton still needs to “choose” which of the DFAs (associated with different states) is used to process the sequence of states that the computation reached at the children of the current node – since the intersection of distinct horizontal languages is empty the choice is unambiguous, however, when beginning to process the sequence of states the automaton has no way of knowing which DFA to use.

A different definition, that we call *strong determinism*, was considered in [5, 15].¹ A strongly deterministic automaton associates to each input symbol a single DFA H_σ equipped with an output function, the state at a parent node labeled by σ is determined (via the output function) by the state H_σ reaches after processing the sequence of states corresponding to the children. Strongly deterministic automata can be minimized efficiently and the minimal automaton is unique [5, 15]. On the other hand, interestingly it was shown in [11] that for weakly deterministic tree automata the minimization problem is NP-complete and the minimal automaton need not be unique.

We study the state complexity of determinizing different variants of nondeterministic tree automata. That is, we develop upper and lower bounds for the size of deterministic tree automata that are equivalent to given nondeterministic automata. We define the size of an unranked tree automaton as a pair of integers consisting of the number of states used in the bottom-up computation, and the sum of the sizes of the NFAs defining the horizontal languages. Note that the two types of states play very different roles in computations of the tree automaton. The other possibility would be, as is done e.g. in [11], to count simply the total number of all states in the different components.

Also, we study the state complexity of the conversions between the strongly and the weakly deterministic tree automata. Although the former model can be viewed to be more restricted, there exist tree languages for which the size of a strongly deterministic automaton is smaller than the size of the minimal weakly deterministic automaton. It turns out to be more difficult to establish lower bounds for the size of weakly deterministic automata than is the case for strongly deterministic automata. Naturally, this can be expected due to the intractability of the minimization of weakly deterministic automata [11].

It should be noted that there are many other deterministic automaton models used for applications on unranked trees, such as stepwise tree automata [3, 4] and nested word automata [1, 14]. Size comparisons between, respectively, stepwise tree automata and strongly deterministic automata or automata operating on binary encodings of unranked trees can be found in [11]. Much work remains to be done on state complexity of tree automata.

To conclude we summarize the contents of the paper. In Section 2 we recall definitions for tree automata operating on unranked trees and introduce some notation. In Section 3 we study the descriptive complexity of conversions between the strongly and the weakly deterministic tree automata, and in Section 4 we study the size blow-up of converting different variants of nondeterministic tree automata to strongly and weakly deterministic automata, respectively. Many of the proofs have been omitted in this extended abstract for the DCFS proceedings.

2 Preliminaries

We assume that the reader is familiar with the basics of formal languages and finite automata [9, 18]. Below we briefly recall some definitions for tree automata operating on unranked trees and fix notations. More details on unranked tree automata and references can be found in [6, 17]. A general reference on tree automata operating on ranked trees is [7].

¹ The paper [5] refers to weak and strong determinism, respectively, as semantic and syntactic determinism.

Basic notions concerning trees, such as the root, a leaf, a subtree, the height of a tree and children of a node are assumed to be known. The set of non-negative integers is \mathbb{N} . A *tree domain* is a prefix-closed subset D of \mathbb{N}^* such that if $ui \in D$, $u \in \mathbb{N}^*$, $i \in \mathbb{N}$ then $uj \in D$ for all $j < i$. The set of nodes of a tree t is represented in the well-known way as a tree domain $\text{dom}(t)$ and the node labeling is given by a mapping $\text{dom}(t) \rightarrow \Sigma$ where Σ is a finite alphabet of symbols. Thus, we use labeled ordered unranked trees. Each node of a tree has a finite number of children with a linear order, but there is no a priori upper bound on the number of children of a node. The set of all Σ -labeled trees is T_Σ .

We introduce the following notation for trees. For $i \geq 0$, $a \in \Sigma$ and $t \in T_\Sigma$, we denote by $a^i(t) = a(a(\dots a(t)\dots))$ a tree, where the nodes $\varepsilon, 1, \dots, 1^{i-1}$ are labelled by a and the subtree at node 1^i is t . When $a \in \Sigma$, $w = b_1 b_2 \dots b_n \in \Sigma^*$, $b_i \in \Sigma$, $1 \leq i \leq n$, we use $a(w)$ to denote the tree $a(b_1, b_2, \dots, b_n)$. When L is a set of strings, $a(L) = \{a(w) \mid w \in L\}$. The set of all Σ -trees where exactly one leaf is labelled by a special symbol x ($x \notin \Sigma$) is $T_\Sigma[x]$. For $t \in T_\Sigma[x]$ and $t' \in T_\Sigma$, $t(x \leftarrow t')$ denotes the tree obtained from t by replacing the unique occurrence of variable x by t' .

A *nondeterministic (unranked) tree automaton* (NTA) is a tuple $A = (Q, \Sigma, \delta, F)$, where Q is the finite set of states, Σ is the alphabet labeling nodes of input trees, $F \subseteq Q$ is the set of final states, and δ is a mapping from $Q \times \Sigma$ to the subsets of Q^* which satisfies the condition that, for each $q \in Q$, $\sigma \in \Sigma$, $\delta(q, \sigma)$ is a regular language. The language $\delta(q, \sigma)$ is called the *horizontal language* associated with q and σ .

A computation of A on a tree $t \in T_\Sigma$ is a mapping $C : \text{dom}(t) \rightarrow Q$ such that for $u \in \text{dom}(t)$, if $u \cdot 1, \dots, u \cdot m$, $m \geq 0$, are the children of u then $C(u \cdot 1) \dots C(u \cdot m) \in \delta(C(u), t(u))$. In case u is a leaf the condition means that $m = 0$ and $\varepsilon \in \delta(C(u), t(u))$.

Intuitively, if a computation of A has reached the children of a σ -labelled node u in a sequence of states q_1, q_2, \dots, q_m , the computation may nondeterministically assign a state q to the node u provided that $q_1 q_2 \dots q_m \in \delta(q, \sigma)$. For $t \in T_\Sigma$, $t^A \subseteq Q$ denotes the set of states that in some bottom-up computation A may reach at the root of t . The *tree language recognized by A* is defined as $L(A) = \{t \in T_\Sigma \mid t^A \cap F \neq \emptyset\}$.

For a tree automaton $A = (Q, \Sigma, \delta, F)$, we denote by $H_{q, \sigma}^A$, $q \in Q$, $\sigma \in \Sigma$, a nondeterministic finite automaton (NFA) on strings recognizing the horizontal language $\delta(q, \sigma)$. The NFA $H_{q, \sigma}^A$ is called a *horizontal automaton*, and states of different horizontal automata are called collectively *horizontal states*. We refer to the states of Q that are used in the bottom-up computation as *vertical states*.

A tree automaton $A = (Q, \Sigma, \delta, F)$ is said to be (semantically) *deterministic* (a DTA) if for $\sigma \in \Sigma$ and any two states $q_1 \neq q_2$, $\delta(q_1, \sigma) \cap \delta(q_2, \sigma) = \emptyset$.

We get a further refinement of classes of automata depending on whether the horizontal languages are defined using DFAs or NFAs. We use $\text{NTA}(M)$ or $\text{DTA}(M)$, respectively, to denote (the class of) nondeterministic or deterministic tree automata where the horizontal languages are specified by the elements in class M . For example, $\text{NTA}(\text{DFA})$ denotes the tree automata where the horizontal languages are recognized by a DFA.

Note that when referring to a tree automaton $A = (Q, \Sigma, \delta, F)$ it is always assumed that the relation δ is specified in terms of automata $H_{q, \sigma}^A$, $q \in Q$, $\sigma \in \Sigma$, and by saying that A is an $\text{NTA}(\text{DFA})$ we indicate that each $H_{q, \sigma}^A$ is a DFA. We refer to $\text{DTA}(\text{DFA})$'s also as *weakly deterministic tree automata* to distinguish them from the below notion of strong determinism.

If A is a $\text{DTA}(\text{NFA})$, for any tree $t \in T_\Sigma$ the bottom-up computation of A assigns a unique vertical state to the root of t , that is, t^A is a singleton set or empty. If the horizontal automata $H_{q, \sigma}^A$ are DFAs, furthermore, for each transition the sequence of horizontal states is processed deterministically. However, as discussed in Section 1, a computation that has reached children of a σ -labeled node in a sequence of states $w \in Q^*$ still needs to make the choice which of the DFAs $H_{q, \sigma}^A$, $q \in Q$, is used to process w . For

this reason we consider also the following notion introduced in [5] that we call strong determinism.

A tree automaton $A = (Q, \Sigma, \delta, F)$ is said to be *strongly deterministic* if for each $\sigma \in \Sigma$, the transitions are defined by a single DFA augmented with an output function as follows. For $\sigma \in \Sigma$ define

$$H_\sigma^A = (S_\sigma, Q, s_\sigma^0, F_\sigma, \gamma_\sigma, \lambda_\sigma), \quad (1)$$

where $(S_\sigma, Q, s_\sigma^0, F_\sigma, \gamma_\sigma)$ is a DFA with set of states S_σ where $s_\sigma^0 \in S_\sigma$ is the start state, $F_\sigma \subseteq S_\sigma$ is the set of final states and $\gamma_\sigma : S_\sigma \times Q \rightarrow S_\sigma$ is the transition function, and λ_σ is a function $F_\sigma \rightarrow Q$. Then we require that for all $q \in Q$ and $\sigma \in \Sigma$: $\delta(q, \sigma) = \{w \in Q^* \mid \lambda_\sigma(\gamma_\sigma(s_\sigma^0, w)) = q\}$. Note that the definition guarantees that $\delta(q_1, \sigma) \cap \delta(q_2, \sigma) = \emptyset$ for any distinct $q_1, q_2 \in Q$, $\sigma \in \Sigma$. The class of strongly deterministic tree automata is denoted as SDTA.²

By the size of an NFA B , denoted $\text{size}(B)$, we mean the number of states of B . Because the roles played by vertical and horizontal states, respectively, in the computations of a tree automaton are essentially different, when measuring the size of an automaton we count the two types of states separately. The size of an NTA(NFA) $A = (Q, \Sigma, \delta, F)$ is defined as

$$\text{size}(A) = [|Q|; \sum_{q \in Q, \sigma \in \Sigma} \text{size}(H_{q, \sigma}^A)] \in (\mathbb{N} \times \mathbb{N}).$$

Using notations of (1), the size of an SDTA A is defined as the pair of integers $\text{size}(A) = [|Q|; \sum_{\sigma \in \Sigma} |S_\sigma|]$.

We make the following notational convention that allows us to use symbols of Σ in the definition of horizontal languages. Unless otherwise mentioned, we assume that a tree automaton always assigns to each leaf symbol labeled σ a state $\bar{\sigma}$ that is not used anywhere else in the computation. That is, for $\sigma \in \Sigma$ and $q \in Q$, $\varepsilon \in \delta(q, \sigma)$ only if $q = \bar{\sigma}$, $\delta(\bar{\sigma}, \sigma) = \{\varepsilon\}$ and $\delta(\bar{\tau}, \sigma) = \emptyset$ for all $\sigma, \tau \in \Sigma$, $\sigma \neq \tau$. When there is no confusion, we denote also $\bar{\sigma}$ simply by σ . When the alphabet Σ is fixed, there is only a constant number of the special states $\bar{\sigma}$ and since, furthermore, the special states have the same function in all types of tree automata, for simplicity, we do not include them when counting the vertical states. The purpose of this convention is to improve readability: many of our constructions become more transparent when alphabet symbols can be used explicitly to define horizontal languages. The convention does not change our state complexity bounds that are generally given within a multiplicative constant.

To conclude this section we give two lemmas that provide lower bound estimates for vertical and horizontal states of SDTAs, respectively. The lower bound condition for vertical states applies, more generally, for DTA(NFA)'s, however, obtaining lower bounds for the number of horizontal states of weakly deterministic automata turns out to be more problematic.

Lemma 2.1 *Let A be an SDTA or a DTA(NFA) with a set of vertical states Q recognizing a tree language L . Assume $R = \{t_1, \dots, t_m\} \subseteq T_\Sigma$ where for any $1 \leq i < j \leq m$ there exists $t \in T_\Sigma[x]$ such that $t(x \leftarrow t_i) \in L$ iff $t(x \leftarrow t_j) \notin L$. Then $|Q| \geq |R| - 1$.*

Lemma 2.2 *Let A be an SDTA with a set of vertical states Q recognizing a tree language L . Let S be a finite set of tuples of Σ -trees and let $b \in \Sigma$. Assume that for any distinct tuples $(r_1, \dots, r_m), (s_1, \dots, s_n) \in S$ there exists $t \in T_\Sigma[x]$ and a sequence of trees u_1, \dots, u_k such that*

$$t(x \leftarrow b(r_1, \dots, r_m, u_1, \dots, u_k)) \in L \text{ iff } t(x \leftarrow b(s_1, \dots, s_n, u_1, \dots, u_k)) \notin L \quad (2)$$

Then the horizontal automaton H_b^A needs at least $|S| - 1$ states.

²Strictly speaking, δ is superfluous in the tuple specifying an SDTA and the original definition of [5] gives instead the automata H_σ^A , $\sigma \in \Sigma$. We use δ in order to make the notation compatible with our other models, and to avoid having to define bottom-up computations of SDTAs separately.

3 Size comparison of the strongly and weakly deterministic tree automata

Here we give upper and lower bounds for the size of a weakly deterministic automaton (a DTA(DFA)) simulating a strongly deterministic one (an SDTA), and vice versa. The computation of a DTA(DFA) can, in some sense, nondeterministically choose which of the horizontal DFAs it uses at each transition. An SDTA does not have this capability and it can be expected that, in the worst case, an SDTA may need considerably more states than an equivalent DTA(DFA). However, there exist also tree languages for which an SDTA can be considerably more succinct than a DTA(DFA).

3.1 Converting an SDTA to a DTA(DFA)

We show that an SDTA can be quadratically smaller than a DTA(DFA). This can be compared with [11] where it was shown that deterministic stepwise tree automata can be quadratically smaller than SDTA's (that are called dPUTA's in [11]).

The upper bound for the conversion is expected but we include a short proof. In the below lemma (and afterwards) we use “ \leq ” to compare pairs of integers componentwise. As introduced in Section 2, for an SDTA A we denote the deterministic automata for the corresponding horizontal languages by H_σ^A , $\sigma \in \Sigma$.

Lemma 3.1 *Let $A = (Q, \Sigma, \delta, F)$ be an arbitrary SDTA.*

We can construct an equivalent DTA(DFA) A' where

$$\text{size}(A') \leq [|Q|; |Q| \times \sum_{\sigma \in \Sigma} \text{size}(H_\sigma^A)]. \quad (3)$$

Proof. For $\sigma \in \Sigma$ denote the components of H_σ^A as in (1). Construct an equivalent DTA(DFA) $A' = (Q, \Sigma, \delta', F)$, where for each $\sigma \in \Sigma$, $q \in Q$, $\delta'(q, \sigma) = \{w \in Q^* \mid \lambda_\sigma(\gamma_\sigma(s_\sigma^0, w)) = q\}$. The languages $\delta'(q_1, \sigma)$ and $\delta'(q_2, \sigma)$, $q_1 \neq q_2$ are always disjoint, and $\delta'(q, \sigma)$ is recognized by a DFA obtained from H_σ^A by choosing as the set of final states $\lambda_\sigma^{-1}(q)$, $q \in Q$, $\sigma \in \Sigma$. The construction does not change the number of vertical states and (3) holds. ■

Next we give a lower bound for the conversion.

Lemma 3.2 *Let $n, z \in \mathbb{N}$ and choose $\Sigma = \{a, b, 0, 1\}$. There exists an SDTA B with input alphabet Σ , n vertical states and $z + 4n$ horizontal states, such that any DTA(DFA) for the tree language $L(B)$ has at least n vertical states and $n(\lfloor \log n \rfloor + 2 + z)$ horizontal states.*

Using Lemma 3.2 with $z = n - \lfloor \log n \rfloor$, we see that the upper bound of Lemma 3.1 is tight within a multiplicative constant. This is stated as:

Theorem 3.1 *An SDTA with n vertical and m horizontal states can be simulated by a DTA(DFA) having n vertical and $n \cdot m$ horizontal states.*

For $n \geq 1$, there exists a tree language L_n recognized by an SDTA with n vertical and $O(n)$ horizontal states such that any DTA(DFA) recognizing L_n has n vertical and $\Omega(n^2)$ horizontal states.

It can be viewed as expected that in the conversion of Theorem 3.1 the number of vertical states does not change. However as will be discussed later, in general, for a DTA(DFA) it may be possible to reduce the number of horizontal states by increasing the number of vertical states.

3.2 Converting a DTA(DFA) to an SDTA

Again we give first an upper bound for the simulation. It is known from ([11] Proposition 24) that the simulation does not increase the number of vertical states.

Lemma 3.3 *Let $B = (Q, \Sigma, \delta, F)$ be an arbitrary DTA(DFA), where $|Q| = n$. Let $H_{q,\sigma}^B = (S_{q,\sigma}, Q, s_{q,\sigma}^0, F_{q,\sigma}, \gamma_{q,\sigma})$ be a DFA for the horizontal language $\delta(q, \sigma)$, $q \in Q$, $\sigma \in \Sigma$.*

We can construct an equivalent SDTA B' where

$$\text{size}(B') \leq [|Q|; \sum_{\sigma \in \Sigma} (\prod_{q \in Q} (|S_{q,\sigma}| - |F_{q,\sigma}|) + \sum_{q \in Q} |F_{q,\sigma}| \cdot \prod_{p \in Q, p \neq q} (|S_{p,\sigma}| - |F_{p,\sigma}|))]$$

If B has m horizontal states, Lemma 3.3 gives for the number of horizontal states of B' a worst-case upper bound that is less than 2^m but is not polynomial in m . Next we give a lower bound construction.

Lemma 3.4 *Let $\Sigma = \{a, b, 0, 1\}$. For any $m \in \mathbb{N}$ and relatively prime numbers $2 \leq k_1 < k_2 < \dots < k_m$, there exists a tree language L over Σ recognized by a DTA(DFA) B with $\text{size}(B) = [m; \sum_{i=1}^m k_i + O(m \log m)]$ such that any SDTA recognizing L has at least m vertical states and $\prod_{i=1}^m k_i$ horizontal states.*

Proof. Let $y_i \in \{0, 1\}^*$ be the binary representation of $i \geq 1$. We define $L = \bigcup_{1 \leq i \leq m} a^i ((b^{k_i})^* y_i)$.

We define for L a DTA(DFA) $B = (Q, \Sigma, \delta, F)$, where $Q = \{q_1, \dots, q_m\}$, $F = \{q_1\}$, $\delta(a, q_i) = (b^{k_i})^* \cdot y_i + q_{i+1}$, for $1 \leq i \leq m-1$, and $\delta(a, q_m) = (b^{k_m})^* \cdot y_m$. Note that the bottom-up computation of B is deterministic because different horizontal languages are marked by distinct binary strings y_i .

Each horizontal language $(b^{k_i})^* \cdot y_i + q_{i+1}$ can be recognized by a DFA with $k_i + \lfloor \log i \rfloor + 3$ states, and in total B has $\sum_{i=1}^m k_i + \sum_{i=1}^m (\lfloor \log i \rfloor) + 3m$ horizontal states (and m vertical states).

Let $B' = (Q', \Sigma, \delta', F')$ be an arbitrary SDTA recognizing L . By choosing $R = \{a(b^{k_i} y_i) \mid 1 \leq i \leq m\} \cup \{a(b)\}$, Lemma 2.1 gives $|Q'| \geq m$.

We show that the DFA $H_a^{B'}$, with notations as in (1), defining transitions corresponding to symbol a needs at least $\prod_{i=1}^m k_i$ states. Suppose that $H_a^{B'}$ has less than $\prod_{i=1}^m k_i$ states. Then there exist $0 \leq j < s < \prod_{i=1}^m k_i$ such that $H_a^{B'}$ reaches the same state after reading strings b^j and b^s , respectively. There must exist $1 \leq r \leq m$ such that k_r does not divide $s - j$. Let $z = j + (k_r - j \bmod k_r)$. Since $H_a^{B'}$ reaches the same state on b^j and b^s , it follows that $H_a^{B'}$ reaches the same state also on $b^z \cdot y_r$ and $b^{z+s-j} \cdot y_r$, respectively. This means that $a^{k_r}(b^z y_r)$ is accepted by B' if and only if $a^{k_r}(b^{z+s-j} y_r)$ is accepted by B' , which is a contradiction because k_r divides z and does not divide $z + s - j$. ■

In the above proof, using a more detailed analysis it could be shown that $H_a^{B'}$ needs $\Omega(m \cdot \log m)$ additional states to process the strings y_i , however, this would not change the worst-case lower bound.

Now we establish that the upper and lower bounds for the DTA(DFA)-to-SDTA conversion are within a multiplicative constant, at least when the sizes of the horizontal DFAs are large compared to the number of vertical states.

Theorem 3.2 *An arbitrary DTA(DFA) $B = (Q, \Sigma, \delta, F)$ has an equivalent SDTA B' with*

$$\text{size}(B') \leq [|Q|; \sum_{\sigma \in \Sigma} \prod_{q \in Q} \text{size}(H_{q,\sigma}^B)], \quad (4)$$

and, for an arbitrary $m \geq 1$ there exists a DTA(DFA) $B = (Q, \Sigma, \delta, F)$ with $|Q| = m$ such that for any equivalent SDTA B' the size of B' has a lower bound within a multiplicative constant of (4).

Proof. The upper bound follows from Lemma 3.3. We get the lower bound from Lemma 3.4 by choosing each k_i to be at least $m \cdot \log m$, $i = 1, \dots, m$. ■

We note that when converting a DTA(DFA) $B = (Q, \Sigma, \delta, F)$ to an equivalent SDTA A , for each $\sigma \in \Sigma$ the horizontal DFA H_σ^A needs at least as many states as a DFA recognizing $L_{B,\sigma} = \bigcup_{q \in Q} \delta(q, \sigma)$. Note that from H_σ^A we obtain a DFA for $L_{B,\sigma}$ simply by ignoring the output function. However, H_σ^A needs to provide more detailed information for a given input string than a DFA simply recognizing $L_{B,\sigma}$, and in fact H_σ^A recognizes the marked union, as formalized below, of the languages $\delta(q, \sigma)$.

We say that a DFA $A = (Q, \Sigma, s_0, F, \gamma)$ equipped with an output function $\lambda : F \rightarrow \{1, \dots, m\}$ recognizes the *marked union* of pairwise disjoint regular languages L_1, \dots, L_m , if $L_i = \{w \in \Sigma^* \mid \lambda(\gamma(s_0, w)) = i\}$, $i = 1, \dots, m$. The following result establishes that the state complexity of marked union may be arbitrarily much larger than the state complexity of union.

Proposition 1 *Let $A = (Q, \Sigma, s_0, F, \gamma, \lambda)$ be a DFA with output function $\lambda : F \rightarrow \{1, \dots, m\}$ that recognizes the marked union of disjoint languages L_i , $i = 1, \dots, m$, and let B be the minimal DFA for $\bigcup_{i=1}^m L_i$.*

Then $\text{size}(A) \geq \text{size}(B)$, and for any $m \geq 1$ there exist disjoint regular languages L_i , $1 \leq i \leq m$, such that $\text{size}(B) = 1$ and the $\text{size}(A) \geq m$.

4 Converting nondeterministic tree automata to deterministic automata

In this section we consider conversions of different variants of nondeterministic automata into equivalent strongly and weakly deterministic automata.

4.1 Converting a nondeterministic automaton to an SDTA

Lemma 4.1 *Let $A = (Q, \Sigma, \delta, F)$ be an NTA(NFA) and for $q \in Q$, $\sigma \in \Sigma$ denote $\text{size}(H_{q,\sigma}^A) = m_{q,\sigma}$.*

(i) *We can construct an equivalent SDTA B where*

$$\text{size}(B) \leq [2^{|Q|}; \sum_{\sigma \in \Sigma} 2^{\left(\sum_{q \in Q} m_{q,\sigma}\right)}]. \quad (5)$$

(ii) *If A is a DTA(NFA), in the upper bound (5) the number of vertical states is at most $|Q|$.*

We do not require the automaton to be complete and, naturally, in (5) the number of vertical states of B could be reduced to $2^{|Q|} - 1$. A similar small improvement could be made to the number of horizontal states, but it would make the formula look rather complicated.

Also, in Lemma 4.1 (ii) the upper bound for the number of horizontal states could be slightly reduced using a more detailed analysis, as in the proof of Lemma 3.3, that takes into account that, in no situation, two distinct NFAs defining the horizontal languages associated with a fixed input symbol σ can accept simultaneously.

Lemma 4.1 did not discuss the case where the bottom-up computation is nondeterministic but the horizontal languages are represented in terms of DFAs. We note that for an NTA(DFA) $A = (Q, \Sigma, \delta, F)$ the construction used in the proof of Lemma 4.1 gives for the size of an equivalent SDTA only the upper bound (5). Although the horizontal languages of A are defined using DFAs, the horizontal languages of the equivalent SDTA B are over the alphabet $\mathcal{P}(Q)$, and this means that the upper bound for the number of horizontal states would not be improved.

Next we state two lower bound results.

Lemma 4.2 *Let $\Sigma = \{a, b\}$. For any relatively prime numbers m_1, m_2, \dots, m_n , there exists a tree language L over Σ such that L is recognized by an NTA(DFA) A with $\text{size}(A) \leq \lceil n; (\sum_{i=1}^n m_i) + 2n - 2 \rceil$, and any SDTA for L needs at least $2^n - 1$ vertical states and $(\prod_{i=1}^n m_i) - 1$ horizontal states.*

Lemma 4.3 *For $n \geq 1$, there exists a tree language L_n recognized by a DTA(NFA) A with n vertical and less than $n \log n$ horizontal states such that for any SDTA B for L_n , $\text{size}(B) \geq \lceil n; 2^n \rceil$.*

The lower bounds given by the above two lemmas are far away from the corresponding upper bounds in Lemma 4.1. Furthermore, we do not have a worst-case construction for general NTA(NFA)'s that would provably give an essentially better lower bound than the one obtained for NTA(DFA)'s in Lemma 4.2.

4.2 Converting a nondeterministic automaton to a DTA(DFA)

We begin with a simulation result establishing an upper bound.

Lemma 4.4 *Let $A = (Q, \Sigma, \delta, F)$ be an NTA(NFA) and for $q \in Q$, $\sigma \in \Sigma$ denote $\text{size}(H_{q,\sigma}^A) = m_{q,\sigma}$.*

(i) *There exists a DTA(DFA) B equivalent to A where*

$$\text{size}(B) \leq \lceil 2^{|Q|}; 2^{|Q|} \cdot (\sum_{\sigma \in \Sigma} 2^{(\sum_{q \in Q} m_{q,\sigma})}) \rceil. \quad (6)$$

(ii) *If A is a DTA(NFA), it has an equivalent DTA(DFA) B where*

$$\text{size}(B) \leq \lceil |Q|; \sum_{q \in Q} \sum_{\sigma \in \Sigma} 2^{m_{q,\sigma}} \rceil.$$

Roughly speaking, the simulation uses a standard subset construction [18] for the set of vertical states, and in order to guarantee that the bottom-up computation remains deterministic the DFA for the horizontal language corresponding to $P \subseteq Q$, $\sigma \in \Sigma$, needs to simulate each horizontal NFA of A corresponding to σ . In the case where A is an NTA(DFA) we do not have a significantly better bound than (6), because the horizontal languages of the DTA(DFA) consist of strings of subsets of Q , which means that we again have to simulate multiple computations of each horizontal DFA of A . In the below lower bound construction of Theorem 4.1 we, in fact, use an NTA(DFA).

We do not have a lower bound that would match the bound of Lemma 4.4. Recall that strongly deterministic automata can be minimized efficiently and the minimal automaton is unique [5], however, minimal DTA(DFA)'s are, in general, not unique and minimization is intractable [11]. When trying to establish lower bounds for the size of a DTA(DFA) $A = (Q, \Sigma, \delta, F)$ there is the difficulty that by adding more vertical states, and hence more horizontal languages, it may still be possible that the total number of horizontal states is reduced. For example, suppose that A has a horizontal language $\delta(q, \sigma) = (a+b)^*b(a+b)^7$, where the minimal DFA has 256 states.³ This language can be represented as a disjoint union of 8 regular languages where the sum of the sizes of the minimal DFAs is only 176. Thus, by replacing the state q by 8 distinct vertical states (that could be equivalent in the bottom-up computation) we could reduce the size of A .

In fact, we do not have a general lower bound condition, analogous to Lemma 2.2, for the number of horizontal states of DTA(DFA)'s and the below lower bound result relies on an ad hoc proof.

³Note that $\delta(q, \sigma)$ is a typical example of a language where the NFA-to-DFA size blow-up is large.

Let $\Sigma = \{a, b\}$. Let p_1, \dots, p_n be the first n primes. Define the tree language

$$T_n = \{a^i(b^k) \mid i \geq 1, k \geq 0, (\exists 1 \leq j \leq n)[k \equiv 0 \pmod{p_j} \text{ and } i \equiv j \pmod{n}]\}.$$

Theorem 4.1 *The tree language T_n can be recognized by an NTA(DFA) A with $\text{size}(A) = \lceil n; (\sum_{i=1}^n p_i) + 2n \rceil$, and for any DTA(DFA) B recognizing T_n ,*

$$\text{size}(B) \geq \lceil 2^n - 1; (2^n - 1) \cdot \prod_{i=1}^n p_i \rceil.$$

Theorem 4.1 gives a construction where converting an NTA(DFA) to a DTA(DFA) causes an exponential blow-up in the number of vertical states, and additionally the size of each of the (exponentially many) horizontal DFAs is considerably larger than the original DFA. However, the size blow-up of the horizontal DFAs does not match the upper bound of Lemma 4.4. In the proof of Theorem 4.1, roughly speaking, we use a particular type of unary horizontal languages in order to be able to (provably) establish that there cannot be a trade-off between the numbers of vertical and horizontal states, and with this type of constructions it seems difficult to approach the worst-case size blow-up of Lemma 4.4.

5 Conclusion

We have studied the state complexity of conversions between different models of tree automata operating on unranked trees. For the conversion of weakly deterministic automata into strongly deterministic automata, and vice versa, we established lower bounds that are within a multiplicative constant of the corresponding upper bound. However, for the size blow-up of converting nondeterministic automata to (strongly and weakly) deterministic automata the upper and lower bounds remain far apart, and this is a topic for further research.

Since a minimal weakly deterministic automaton need not be unique [11], it is, in general, hard to establish lower bounds for the number of horizontal states of weakly deterministic automata and we do not have tools like Lemma 2.2 that is used for strongly deterministic automata. Weakly deterministic automata can have trade-offs between the numbers of vertical and horizontal states, respectively, and it would be useful to establish some upper bounds for how much the number of horizontal states can be reduced by introducing additional vertical states.

References

- [1] J. Alur, P. Madhusudan. Adding nesting structure to words. *J. Assoc. Comput. Mach.* 56(3), 2009.
- [2] A. Brüggemann-Klein, M. Murata, D. Wood. Regular tree and regular hedge languages over unranked alphabets. HKUST Technical report, 2001.
- [3] J. Carme, J. Niehren, M. Tommasi. Querying unranked trees with stepwise tree automata. In *Proceedings of RTA'04, Lect. Notes Comput. Sci.*, 3091, 105–118, Springer, 2004.
- [4] J. Champavère, R. Gilleron, A. Lemay, J. Niehren. Efficient inclusion checking for deterministic tree automata and XML schemas. *Inform. Comput.* 207, 1181–1208, 2009.
- [5] J. Cristau, C. Löding, W. Thomas. Deterministic automata on unranked trees. In *Proc. of FCT'05, Lect. Notes Comput. Sci.*, 3623, 68–79, Springer, 2005.
- [6] H. Comon, M. Dauchet, R. Gilleron, F. Jacquemard, D. Lugiez, C. Löding, S. Tison, M. Tommasi. *Tree Automata Techniques and Applications*, electronic book available at tata.gforge.inria.fr, 2007.

- [7] F. Gécseg, M. Steinby. Tree languages. In G. Rozenberg, A. Salomaa (eds.), *Handbook of Formal Languages*, vol. III, 1–68, Springer, 1997.
- [8] M. Holzer, M. Kutrib. Descriptive and computational complexity of finite automata. *Proc. of LATA'09, Lect. Notes Comput. Sci.*, 5457, 23–42, Springer, 2009.
- [9] J.E. Hopcroft, J.D. Ullman. *Introduction to Automata Theory, Languages, and Computation*, Addison Wesley, 1979.
- [10] G. Liu, C. Martin-Vide, A. Salomaa, S. Yu. State complexity of basic language operations combined with reversal. *Inform. Comput.* 206, 1178–1186, 2008.
- [11] W. Martens, J. Niehren. On the minimization of XML schemas and tree automata for unranked trees. *J. Comput. System Sci.* 73, 550–583, 2007.
- [12] T. Milo, D. Suci, V. Vianu. Typechecking for XML transformers. *J. Comput. System Sci.* 66, 66–97, 2003.
- [13] F.R. Moore. On the bounds for state-set size in the proofs of equivalence between deterministic, nondeterministic, and two-way finite automata. *IEEE Transactions on Computers* 20, 1211–1214, 1971.
- [14] X. Piao, K. Salomaa. Operational state complexity of nested word automata. *Theoret. Comput. Sci.* 410, 3290–3302, 2009.
- [15] S. Raeymaekers, M. Bruynooghe. Minimization of finite unranked tree automata. Manuscript, 2004.
- [16] K. Salomaa. Descriptive complexity of nondeterministic finite automata. In *Proc. DLT'07, Lect. Notes Comput. Sci.* 4588, 31–35, Springer, 2007.
- [17] T. Schwentick. Automata for XML, *J. Comput. System Sci.* 73, 289–315, 2007.
- [18] S. Yu. Regular languages. In G. Rozenberg, A. Salomaa (eds.), *Handbook of Formal Languages*, vol. I, . 41–110, Springer, 1997.